

### **REMARKS / ARGUMENTS**

Claims 18-21, 25 and 33-41 are pending in the instant application. Claims 3 and 22-24 have been previously cancelled, and claims 1-2, 4-17, 26-32 and 42-49 have been withdrawn due to allegedly being directed to a non-elected invention. Claims 18, 33, 36 and 39 are independent. Claims 19-21, 25, 34-35, 37-38 and 40-41 depend directly or indirectly from independent claims 18, 33, 36 and 39, respectively.

Claims 18, 20, 21, 25 and 36-41 are rejected under 35 USC 103(a) as being unpatentable over USPP 2001/0037406 ("Philbrick"), in view of USPP 2002/0059451 ("Haviv") and USPP 2003/0046330 ("Hayes"). Claims 33-35, 30 are rejected under 35 USC 103(a) as being unpatentable over Philbrick in view of Microsoft Winsock Direct and Protocol Offload on SANs, 03/03/2001 ("Microsoft") and Hayes. Claim 19 is rejected under 35 USC 103(a) as being unpatentable over Philbrick in view of Haviv and Hayes.

#### **I. Examiner's Response to Arguments in the Office Action**

The Applicant previously filed an Appeal Brief on 4/19/2010. The Examiner reopened prosecution and raised a new ground of rejection based on combining Philbrick and Haviv with a new reference Hayes. Regarding independent claim 18, the Examiner concedes the following regarding Philbrick and Haviv (see the Office Action in pages 4-5):

"Philbrick discloses a network interface card (NIC) with 4 different connectors for the purpose of serving 4 different conduits (fig. 16 and [0066], [0067] and [0106]). Philbrick-Haviv does not disclose a network interface card with a single Ethernet connector."

The Examiner looks to Hayes to overcome Philbrick and Haviv's above deficiencies, and states the following (see the Office Action in page 5):

However, Hayes discloses a NIC with a single connector (fig. 5, [0026], single PHY interface connected to the network, [0018], [0020], NIC handling both offload protocol traffics and regular traffics)"

The Examiner relies for support on Hayes' Fig. 5, and equates Hayes' PHY interface 18 connected to the computer network E to Applicant's "single Ethernet connector". The Applicant respectfully disagrees, and points out that Hayes' Fig. 5 clearly discloses that the PHY interface 18 is part of the NIC (the alleged "single integrated convergent network controller (ICNC) chip"). In this regard, Hayes does not disclose or suggest the alleged "a single Ethernet connector for handling a plurality of different types of network traffic..." or "the single Ethernet controller chip is coupled to the single ICNC chip", and Hayes still does not overcome the deficiencies of Philbrick and Haviv.

Assuming arguendo, that Hayes' PHY interface 18 is the alleged "single Ethernet connector" (which it is not), the Examiner's argument is still deficient. For example, the Examiner seems to equate Hayes' offload protocol traffic and regular traffic as Applicant's "plurality of different types of network traffic". The Applicant respectfully disagrees, and points out that Hayes' offload protocol traffic and regular traffic are of the

same traffic type, except through different processing paths (much like Philbrick's), namely, via the fast path which by-passes the host protocol stack (i.e., via the Auxiliary processor AP), or via the regular path where the packets are processed by the host protocol stack in the host operating system OS.

For example, the Examiner is referred to the following citations of Hayes:

**"[0018] In another preferred embodiment of the invention, the auxiliary processor offloads the transmission of iSCSI data over the TCP/IP network protocol, performing all necessary TCP/IP functions that occur during the normal course of a TCP/IP transmit operation and all necessary iSCSI protocol functions. In the event of an error or other exceptional condition, the auxiliary processor transfers control back to the offloading host to handle the condition.**

[0020] In other preferred embodiments, other network protocols, transport protocols and application protocols may be offloaded to the auxiliary processor. **The protocol may be a combination of protocols including the network protocol, the transport protocol and the application protocol.** The offloaded protocols can be any protocol or set of protocols in the seven layer ISO protocol reference model. When multiple protocols of different layers are taken together, each unique combination of protocols is treated as a separate protocol. This capability allows the underlying protocols to be tailored to the requirements of the application and the application protocol. The additional protocols are described in detail below."

See Hayes at paragraphs [0018] and [0020] (emphasis added). Hayes' paragraph [0020] discloses that "each unique combination of protocols (i.e., including the network protocol, the transport protocol and the application protocol) is treated as a separate protocol. Hayes' paragraph [0018] provides an example of offloading such a combination, such as the iSCSI protocol functions (i.e., the application protocol) over the

TCP/IP (i.e., the network and transport protocol) via the auxiliary processor (i.e., by-passing the host protocol stack in the OS). Hayes' paragraph [0018] further discloses that for the same iSCSI protocol functions (i.e., the application protocol) over the TCP/IP (i.e., the network and transport protocol), the auxiliary processor transfers control back to the offloading host to handle the condition "in the event of an error or other exceptional condition".

In other words, Hayes discloses that the same iSCSI protocol function can be processed via the fast path (i.e., offloading via the auxiliary processor" in the NIC) or the regular path (i.e., the host protocol stack processing path), the protocol function being the same "network traffic type" of iSCSI protocol.

Based on the foregoing rationale, the Applicant maintains that Hayes' PHY interface 18 (i.e., the alleged "single Ethernet connector") does not handle the alleged "plurality of different types of network traffic". In this regard, Hayes does not disclose or suggest "a single Ethernet connector for handling a plurality of different types of network traffic...", or "the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic," as recited in Applicant's claim 18. Therefore, Hayes does not overcome the above deficiencies of Philbrick and Haviv. Accordingly, Applicant's claim 18 is submitted to be allowable.

## II. REJECTION UNDER 35 U.S.C. § 103

In order for a *prima facie* case of obviousness to be established, the Manual of Patent Examining Procedure ("MPEP") states the following:

"First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or combine the teaching. Second, there must be a reasonable expectation of success. Finally, **the prior art reference (or references when combined) must teach or suggest all the claim limitations.** The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure."

See MPEP at § 2142, citing *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991) (emphasis added). Further, MPEP § 2143.01 states that "the mere fact that references can be combined or modified does not render the resultant combination obvious unless the prior art suggests the desirability of the combination," and that "although a prior art device 'may be capable of being modified to run the way the apparatus is claimed, there must be a *suggestion or motivation in the reference to do so*'" (citing *In re Mills*, 916 F.2d 680, 16 USPQ 2d 1430 (Fed. Cir. 1990)). Moreover, MPEP § 2143.01 also states that the level of ordinary skill in the art cannot be relied upon to provide the suggestion..., citing *Al-Site Corp. v. VSI Int'l Inc.*, 174 F.3d 1308, 50 USPQ 2d 1161 (Fed. Cir. 1999). Additionally, if a *prima facie* case of obviousness is not established, the Applicant is under no obligation to submit evidence of nonobviousness.

The examiner bears the initial burden of factually supporting any *prima facie* conclusion of obviousness. If the examiner does not produce a *prima facie* case, the applicant is under no obligation to submit evidence of nonobviousness.

See MPEP at § 2142.

**A. The Proposed Combination of Philbrick, Haviv and Hayes Does Not Render Claims 18, 20, 21, 25 and 36-41 Unpatentable**

The Applicant turns to the rejection of claims 18, 20, 21, 25 and 36-41 under 35 U.S.C. § 103(a) as being unpatentable over Philbrick in view of Haviv and Hayes.

**A(1) Independent Claims 18, 36 and 39**

With regard to the rejection of independent claim 18 under 35 U.S.C. § 103(a), the Applicant submits that Philbrick does not disclose or suggest “a single Ethernet connector for handling a plurality of different types of network traffic transported via a single fabric, wherein the single Ethernet connector is coupled to the single integrated convergent network controller chip,” or “the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic,” as recited in Applicant’s claim 18. The Examiner states the following:

“For claim 18, Philbrick discloses a server, comprising: a **single integrated convergence network controller chip** (fig. 6, fig. 1, **network interface card INIC 22**);

an Ethernet connector for handling **a plurality of different types of network traffic** (fig. 16, one of the Ethernet connectors for receiving multiple traffic types, [0065], SCSI and TCP, or Etherstorage or SEP and TCP, [0069] lines 20-23, different storage protocols over TCP/IP, [0084], [0085], NAS traffic and network storage traffic over network line 644, utilizing iSCSI and TCP/NetBios/SMB, [0085], iSCSI and TCP/NetBios/SMB, fig. 15, [0093], [0097], [0099], fast path audio and video traffics and real time voice/video traffics and NAS, RTP/RTCP and SIP and MGCP)),

the single connector is coupled to the single integrated convergent network controller chip ([0066] lines 12-15, Ethernet connector 424 coupled to the INIC)

the single integrated convergence network controller chip is operable to concurrently process the plurality of different types of traffic ([0065] lines 15-21, at least two traffics SCSI and TCP/IP, fig. 14, [0084], [0085], NAS traffic and network storage traffic over network line 644, utilizing iSCSI and TCP/NetBios/SMB, [0085], iSCSI and TCP/NetBios/SMB, fig. 15, [0093], [0097], [0099], fast path audio and video traffics and real time voice/video traffics and NAS, RTP/RTCP and SIP and MGCP))."

See Office Action at pages 3-4 (emphasis added). The Examiner is referred to the Applicant's arguments in the 4/19/2010 Brief on Appeal. Namely, Philbrick does not disclose or suggest "the single integrated convergent network controller chip is operable to **concurrently process** the plurality of different types of network traffic (via the single fabric and a single Ethernet connector) **for the plurality of servers,**" or "... **the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic** for the plurality of servers, which is transported via the single fabric," as recited in Applicant's claim 18.

. The Examiner seemed to allege that Philbrick's Fig. 14 and ¶¶0084-0085 disclose that the INIC 622 (the alleged "single integrated convergent network controller chip") handles at least two types of traffic over the same fabric (i.e., network line 644), namely, the iSCSI SAN traffic and TCP/NetBios/SMB NAS traffic. The Applicant respectfully disagrees. Even though Philbrick discloses that the same network line 644 (the alleged "single fabric") is connected to both the SAN storage unit 640 and the NAS storage unit 642, Philbrick discloses that **only one storage unit is accessed at a time**

(i.e., **there is no concurrent access**). In other words, there is still only one type of traffic, **either** the iSCSI SAN traffic or the TCP/NetBios/SMB NAS traffic (but **not both**) are handled by the network line 644 (the alleged "single fabric"). The Examiner is referred to the following citation of Philbrick:

"A storage fast-path is provided by the INIC 622, under control of the server, for **data transferred between network storage units 640 or 642** and client 602 that does not cross the I/O bus. Data is communicated between INIC 622 and network storage unit 640 in accordance with a block format, such as SCSI/TCP or iSCSI, whereas data is communicated between INIC 622 and NAS storage unit 642 in accordance with a file format, such as TCP/NetBios/SMB. For either storage fast-path **the INIC 622 may hold another CCB defining a connection with storage unit 640 or 642.**"

See Philbrick at ¶0085 (emphasis added). As seen in the above citation, Philbrick clearly discloses that **either** the iSCSI data or the TCP/NetBios/SMB data is transferred to the INIC 622. In this regard, **the network line 644** (the alleged "single fabric") still **handles one type of traffic at a time, but not concurrently**.

In addition, the Examiner had argued before, that the iSCSI traffic includes two separate protocol types, namely, the SCSI and TCP network protocols, and therefore, Philbrick's INIC allegedly performs "**concurrently processing the plurality of different types of network traffic**". The Applicant respectfully disagrees, and refers the Examiner to Applicant's arguments in the 7/23/09 response (*see* pages 16-17), that the iSCSI traffic is a **single traffic type**, not two traffic types. Specifically, the Examiner is referred to Satran, which discloses that the **SCSI** (not iSCSI) protocol and commands



by themselves cannot be transported across the TCP/IP network infrastructure. In this regard, Satran clearly discloses that the **iSCSI protocol is a stand alone new protocol standard, with its own commands and numbering scheme for the iSCSI PDUs, which are transported over the TCP/IP network infrastructure.**

Therefore, contrary to the Examiner's allegation, the iSCSI protocol by itself, is a unique protocol type (i.e., a single protocol of a single traffic type), and is not considered as two separate protocol types (i.e., separate SCSI and TCP network traffics). In this regard, Philbrick's INIC handles the iSCSI traffic as a **single traffic type only**, and not two traffic types, as alleged by the Examiner. Accordingly, Philbrick's INIC does not disclose or suggest "concurrently processing the plurality of different types of network traffic," as recited in Applicant's claim 18.

The Examiner further argued the following in the 10/6/2009 Final Office Action:

"In response to arguments of Philbrick's INIC does not process a plurality of traffics concurrently. In traversal, Philbrick clearly discloses that the **server INIC holds 3 CCB's concurrently for distinguishing 3 different types of network traffics**, therefore supporting a chip concurrently process a plurality of traffic types ([0085], 3 CCBs for different traffic types so that the server INIC can process the traffic types according to the server protocol stack in fig. 14)."

See 10/6/2009 Final Office Action in page 4 (emphasis added). The Examiner alleged that Philbrick's Fig. 14 and ¶0085 disclose that the INIC 622 allegedly concurrently holds three CCBs (i.e., client CCB, SAN CCB and NAS CCB) for distinguishing three different types of network traffic (i.e., TCP, iSCSI and SMB),

therefore allegedly supporting a chip that concurrently processes a plurality of traffic types. The Applicant points out that the Examiner seems to have misinterpreted Philbrick's ¶0085, consequently, the Examiner's above allegation is contrary to Philbrick's disclosure. The Examiner is respectfully referred the above citation of Philbrick in ¶0085:

"A storage fast-path is provided by the INIC 622, under control of the server, for **data transferred between network storage units 640 or 642** ... For either storage fast-path **the INIC 622 may hold another CCB defining a connection with storage unit 640 or 642.**"

As seen, Philbrick clearly discloses that the INIC 622 controls data transfer between **SAN** network storage unit 640 or the **NAS** storage unit 642. In addition, Philbrick also clearly discloses that the INIC 622 holds only one CCB connection at a time, namely, a NAS CCB for the NAS storage unit 640, or a SAN CCB for the SAN storage unit 642.

In this regard, contrary to the Examiner's allegation, Philbrick does not disclose or suggest that the SAN CCB and the NAS CCB are concurrently handled by the INIC 622. Moreover, the Applicant also points out that Philbrick's client CCB is handled on a separate network line (not in the same alleged "single fabric") using a separate connector. Therefore, the client CCB is not even handled by the same alleged "single connector" or the alleged "single fabric". Accordingly, Philbrick's INIC 622 does not process all three CCBs (i.e., client CCB, SAN CCB and NAS CCB) concurrently (in a single fabric and using a single Ethernet connector).

Based on the foregoing rationale, the Applicant maintains that Philbrick does not disclose or suggest "... the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic for the plurality of servers, which is transported via the single fabric," as recited in Applicant's claim 18. Haviv does not overcome Philbrick's above deficiencies.

The Examiner in the 10/6/2009 Final Office Action (see page 6) further alleges that Philbrick (see Fig. 15 and ¶¶0093, 0097, 0099) discloses other traffic, such as the fast path audio and video traffic, as well as real time voice/video and NAS RTP/RTCP, SIP and MGCP that are communicated to the client server 602, allegedly **via a single fabric**.

The Applicant respectfully disagrees, and points out that the listed traffics (i.e., fast path audio and video traffic, as well as real time voice/video and NAS RTP/RTCP, SIP and MGCP) are **not** communicated to the client server 602 **via** an alleged "**fabric**". Instead, Philbrick (see Fig. 15 and ¶¶0093, 0097, 0099) discloses the traffic is communicated **via an I/O bus 675**, which is neither coupled to an Ethernet connector nor to the INIC 606 of the client server 602.

In this regard, Philbrick's (see Fig. 15 and ¶¶0093, 0097, 0099) traffic in the I/O bus 675 does not read on "a single Ethernet connector for handling a plurality of different types of network traffic transported via a single fabric fabric, wherein the **single** Ethernet connector is coupled to the single integrated convergent network controller

chip ...," as recited in Applicant's claim 18. Haviv does not overcome Philbrick's above deficiencies.

Based on the foregoing rationale, the Applicant maintains that Philbrick does not disclose or suggest "the single integrated convergent network controller chip is operable to **concurrently process** the plurality of different types of network traffic (via the single fabric and a single Ethernet connector) **for the plurality of servers,**" or "... **the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic** for the plurality of servers, which is transported via the single fabric," as recited in Applicant's claim 18.

The Examiner concedes the following regarding Philbrick and Haviv (see the Office Action in pages 4-5):

"Philbrick discloses a network interface card (NIC) with 4 different connectors for the purpose of serving 4 different conduits (fig. 16 and [0066], [0067] and [0106]). Philbrick-Haviv does not disclose a network interface card with a single Ethernet connector."

The Examiner looks to Hayes to overcome Philbrick and Haviv's above deficiencies, and states the following (see the Office Action in page 5):

However, Hayes discloses a NIC with a single connector (fig. 5, [0026], single PHY interface connected to the network, [0018], [0020], NIC handling both offload protocol traffics and regular traffics)"

The Examiner is referred to Applicant's above arguments in section I, namely, Hayes' Fig. 5 clearly discloses that the PHY interface 18 is part of the NIC (the alleged "single integrated convergent network controller (ICNC) chip"). In this regard, Hayes

does not disclose or suggest the alleged “a single Ethernet connector for handling a plurality of different types of network traffic...” or “the single Ethernet controller chip is coupled to the single ICNC chip”, and Hayes still does not overcome the deficiencies of Philbrick and Haviv.

Assuming arguendo, that Hayes’ PHY interface 18 is the alleged “single Ethernet connector” (which it is not), the Examiner’s argument is still deficient. For example, the Examiner seems to equate Hayes’ offload protocol traffics and regular traffics as Applicant’s “plurality of different types of network traffic”. The Applicant respectfully disagrees, and points out that Hayes’ offload protocol traffics and regular traffics are of the same traffic type, except through different processing paths (much like Philbrick’s), namely, via the fast path which by-passes the host protocol stack (i.e., via the Auxiliary processor AP), or via the regular path where the packets are processed by the host protocol stack in the host operating system OS.

For example, Hayes’ paragraph [0020] discloses that “each unique combination of protocols (i.e., including the network protocol, the transport protocol and the application protocol) is treated as a separate protocol. Hayes’ paragraph [0018] provides an example of offloading such a combination, such as the iSCSI protocol functions (i.e., the application protocol) over the TCP/IP (i.e., the network and transport protocol) via the auxiliary processor (i.e., by-passing the host protocol stack in the OS). Hayes’ paragraph [0018] further discloses that for the same iSCSI protocol functions (i.e., the application protocol) over the TCP/IP (i.e., the network and transport protocol),

the auxiliary processor transfers control back to the offloading host to handle the condition "In the event of an error or other exceptional condition".

In other words, Hayes discloses that the same iSCSI protocol function can be processed via the fast path (i.e., offloading via the auxiliary processor" in the NIC) or the regular path (i.e., the host protocol stack processing path), being the same "network traffic type" of iSCSI protocol.

Based on the foregoing rationale, the Applicant maintains that Hayes' PHY interface 18 (i.e., the alleged "single Ethernet connector") does not handle the alleged "plurality of different types of network traffic". In this regard, Hayes does not disclose or suggest "a single Ethernet connector for handling a plurality of different types of network traffic...", or "the single integrated convergent network controller chip is operable to concurrently process the plurality of different types of network traffic," as recited in Applicant's claim 18. Therefore, Hayes does not overcome the above deficiencies of Philbrick and Haviv. Accordingly, Applicant's claim 18 is submitted to be allowable.

Likewise, independent claims 36 and 39 are similar in many respects to claim 18, are also submitted to be allowable based on the same rationale of claim 18.

**B. Dependent Claims 20-21, 25 and 37-38 and 40-41**

Based on at least the foregoing, the Applicant believes the rejection of independent claims 18, 36 and 39 under 35 U.S.C. § 103(a) as being unpatentable by

the combination of Philbrick, Haviv and Hayes has been overcome and requests that the rejection be withdrawn. Additionally, claims 20, 21, 25 and 37-38 and 40-41 depend directly or indirectly from independent claims 18, 36 and 39 and are, also respectfully submitted to be allowable.

**B(1). Rejection of Dependent Claims 25, 38 and 41**

The Examiner states the following in the Office Action:

"For claim 25, Philbrick-Haviv-Hayes further discloses the plurality of different types of traffic comprises at least two of network traffic, storage traffic, interprocess communication (IPC) traffic and cluster traffic (Philbrick, [0065] lines 15-21, network traffic TCP/IP and storage traffic SCSI, Haviv, [0019])."

See Office Action in page 5. The Examiner is referred to Applicant's above arguments in claim 18, namely, that **the iSCSI traffic is a single type of traffic comprising iSCSI PDU messages**. Also, Philbrick's Fig. 14 and ¶0085 disclose that only one type of network traffic, **either the SAN or NAS traffic**, is handled by the NIC 622 (the alleged "integrated chip"). In this regard, Philbrick does not disclose that "at least two of network traffic, storage traffic, IPC traffic or the cluster traffic", are **concurrently** handled by the Ethernet connector and the integrated chip. Claim 25 is therefore allowable. Claims 38 and 41 are allowable for the same rationale as stated with regard to claim 25.

**B(2). Rejection of Dependent Claims 37 and 40**

The Examiner states the following in the Office Action:

"For claim 37, Philbrick-Haviv-Hayes discloses said single integrated convergent network controller chip comprises a layer 2 network interface card (L2 NIC) (Philbrick, [0065] lines 7-11, Ethernet, fig. 24, MAC controller), a transmission control protocol (TCP) processor, an iSCSI processor ([0065] lines 15-21, iSCSI processing over TCP/IP) and a remote direct memory access (RDMA) processor (fig. 25, DMA controller), and a Management Agent processor ([0106], last sentence)."

See Office Action in page 8. The Examiner relies for support on Philbrick's Fig. 6 and paragraph [0065], and alleges that Philbrick's INIC 400 discloses a TCP processor and an iSCSI processor. The Applicant respectfully disagrees, and refers the Examiner to the following citation of Philbrick:

**"SANS 418 and 420 may run a storage protocol such as SCSI over TCP/IP** or SCSI Encapsulation Protocol. One such storage protocol is ... "iSCSI (Internet SCSI) June 2000, which in an earlier Internet-Draft was entitled "SCSI/TCP (SCSI over TCP)," . employs SCSI Encapsulation Protocol (SEP) at the session layer, and either TCP or SAN transport protocol (STP) at the transport layer, depending primarily upon whether data is being transferred over a WAN or the Internet, for which TCP is used, or **data is being transferred over a LAN or SAN**, for which STP is used."

See Philbrick in ¶[0065] (emphasis added). The Applicant points out that Philbrick, in the above citation, discloses that it is the SANS 418 and 420, not the NIC 400 in the host server, which run the iSCSI protocol (i.e., SCSI over TCP). In this regard, Philbrick's Fig. 6 and ¶[0065] do not disclose a TCP processor or iSCSI



processor. On the contrary, Philbrick's Fig. 14 discloses that the TCP protocol stack is processed in the host, and the SCSI protocol is processed in the storage unit 634.

In addition, Philbrick's Fig. 25 merely discloses a DMA controller 2206, which performs DMA operations. However, Philbrick's Fig. 25, as well as the entire reference, does not disclose RDMA protocol processing. In this regard, Philbrick does not disclose or suggest the alleged "RDMA protocol processor" in the INIC, as alleged by the Examiner.

Likewise, the Examiner alleges that Philbrick's ¶0106 discloses that the INIC includes a management agent processor. The Applicant respectfully disagrees, and points out that Philbrick's ¶0106 states:

"The MAC units 722, 724, 726 and 728 also provide statistical information that can be used for simple network management protocol (SNMP)."

Philbrick merely discloses that the MAC units **provide statistical information that can be used for SNMP**. There is no disclosure or suggestion that the SNMP protocol is actually processed within the INIC. In this regard, the Examiner's allegation that Philbrick discloses a management agent processor in the INIC, is unsupported.

Based on the foregoing rationale, the Applicant maintains that Philbrick and Hayes at least do not disclose or suggest "wherein said single integrated convergent network controller chip comprises a layer 2 network interface card (L2 NIC), a **transmission control protocol (TCP) processor, an iSCSI processor, a remote**

**direct memory access (RDMA) processor and a Management Agent processor,"** as recited in claim 37. Haviv does not overcome Philbrick and Hayes' above deficiencies. Claim 37 is therefore allowable. Claim 40 is allowable for the same rationale as stated with regard to claim 37.

**III. The Proposed Combination over Philbrick, Microsoft and Hayes Does Not Render Claims 33-35 Unpatentable**

The Applicant turns to the rejection of claims 33-35 under 35 U.S.C. § 103(a) as being unpatentable over Philbrick and further in view of Microsoft and Hayes.

**A. Independent Claim 33**

With regard to the rejection of independent claim 33 under 35 U.S.C. § 103(a), the Applicant refers the Examiner to the arguments in claim 1, namely, A) Philbrick and Hayes do not disclose or suggest "a single integrated convergent network controller chip that is enabled to **concurrently process a plurality of different types of traffic**"; and B) Philbrick and Hayes do not disclose or suggest "**a single connector coupled to a single integrated convergent chip** that is enabled to concurrently process a plurality of different types of traffic".

Microsoft does not overcome the deficiencies of Philbrick and Hayes. Accordingly, a prima facie case of obviousness cannot be established by the combination of Philbrick, Microsoft and Hayes to reject claim 33, therefore, claim 33

should be allowable. The Applicant respectfully requests that the rejection of claim 33 under 35 U.S.C. § 103(a) be withdrawn.

**B. Rejection of Dependent Claims 34-35**

Based on at least the foregoing, the Applicant believes the rejection of the independent claim 33 has been overcome. Additionally, claims 34-35 depend from independent claim 33, and are, also respectfully submitted to be allowable.

**IV. The Proposed Combination over Philbrick, Haviv and Hayes and what was known in the art Does Not Render Claim 19 Unpatentable**

Based on at least the foregoing, the Applicant believes the rejection of the independent claim 18 has been overcome. Additionally, claim 19 depends from independent claim 18, and is, also respectfully submitted to be allowable.

The Applicant reserves the right to argue additional reasons beyond those set forth herein to support the allowability of dependent claims 18-21, 25 and 33-41, should such a need arise.

**CONCLUSION**

Based on at least the foregoing, the Applicant believes that all pending claims 18-21, 25 and 33-41 are in condition for allowance. If the Examiner disagrees, the Applicant respectfully requests a telephone interview, and requests that the Examiner telephone the undersigned Patent Agent at (312) 775-8093.

The Commissioner is hereby authorized to charge any additional fees or credit any overpayment to the deposit account of McAndrews, Held & Malloy, Ltd., Account No. 13-0017.

A Notice of Allowability is courteously solicited.

Respectfully submitted,

Date: September 20, 2010

/ Frankie W. Wong /

Frankie W. Wong  
Registration No. 61,832  
Patent Agent for Applicant

MCANDREWS, HELD & MALLOY, LTD.  
500 WEST MADISON STREET, 34TH FLOOR  
CHICAGO, ILLINOIS 60661  
(312) 775-8093 (FWW)  
Facsimile: (312) 775 – 8100